Policy Rules and Targets: Framing the Central Banker's Problem

Stephen G. Cecchetti

Central bank policymakers are not primarily random number generators.[1] Reading both the financial press and the work of academics, however, one might get the opposite impression. Reporters (and the readers of their stories) seem to attach considerable importance to each Federal Open Market Committee policy decision. Academic work on the impact of central bank policy gives a similar impression, as statistical procedures produce a time series of pure white noise innovations that are labeled "policy shocks."[2] But central bankers expend substantial energy attempting to tailor their actions to current economic conditions. In other words, policymakers are reacting to the environment, not injecting noise.

But what is central bank policy anyway? The policymaker's problem can be characterized in the following way. Using an instrument such as an interest rate, together with knowledge of the evolution of the economy (aggregate output and the price level), the policymaker seeks to stabilize output and prices about some path that is thought to be optimal. In carrying out this goal, the policymaker must often trade off variability in output for variability in prices because it is generally not possible to stabilize both. This process yields what most people would call a policy rule, that is, a systematic rule for adjusting the quantity that the central bank controls as the state of the economy fluctuates. In other words, the study of policy should focus on the systematic portion of policymakers' actions, not the shocks.

In this essay, I discuss a number of conceptual and practical issues associated with viewing policymaking in this analytical framework. These issues include the implications for policymaking of the slope of the output-inflation variability trade-off, the influence of various types of uncertainty on the policymaker's problem, the consequences of the fact that the nominal interest rate cannot fall below zero, and possible justifications for interest rate smoothing.

Stephen G. Cecchetti is executive vice president and director of research at the Federal Reserve Bank of New York.

Although my intention is to raise, rather than resolve, key questions concerning the formulation of a policy rule, I do offer important new evidence on one point. This concerns the potential consequences of the move by many central banks to adopt some form of price-level or inflation targeting. In taking this approach, central banks are implicitly altering the relative importance of inflation and output variability in their objectives, increasing the weight they attach to the former relative to the latter. But the data suggest that the output-inflation variability trade-off is extremely steep, implying that an effort to decrease inflation variability modestly could lead to a significant increase in output variability. Thus, policymakers considering pure inflation targeting should be aware that their change in emphasis could have undesirable side effects.

An Analytical Framework for Policy Formulation

As I suggested in the introduction, central bank policy can be thought of as the solution to a problem in which the policymaker uses an interest rate to stabilize the variability of output and prices about some path. A truly complete

Central bank policy can be thought of as the solution to a problem in which the policymaker uses an interest rate to stabilize the variability of output and prices about some path.

description of the policymaker's problem begins with an intertemporal general equilibrium model based on a social welfare function (tastes), production functions (technology), and market imperfections that cause nominal shocks to have real effects (nominal rigidities). The goal would be welfare maximization.

I do not propose to delineate the fully specified problem. Instead, I begin with a commonly used quadratic loss function that might be a second-order approximation to the objective function in this more detailed problem.[3] The policymaker seeks to minimize the discounted sum of squared deviations of output and prices from their target paths. The general form of such a loss function can be written as

$$L = E_t\left(\sum_{i=0}^{h}\beta^i\{\alpha[p_{t+i}-p^*_{t+i}]^2 + (1-\alpha)[y_{t+i}-y^*_{t+i}]^2\}\right), \quad (1)$$

where p_t is the (log) aggregate price level, y_t is the (log) aggregate output, p^* and y^* are the desired levels for p and y, β is the discount factor, h is the horizon, α is the relative weight given to squared price and output deviations from their desired paths, and E_t is the expectation conditional on information at time t.[4] The loss function provides the policymaker with information about preferences over different paths for the variance of output and prices.

A complete formulation of L requires description of p^* and y^*. I will focus on the desired price path, ignoring issues concerning y^*.[5] Here we encounter the following question: Should the objective be a price-level path or an inflation rate? The first of these, *level targeting*, would dictate that

$$p^*_t = p^*_{t-1} + \pi^* = \pi^* t, \quad (2)$$

where π^* is the desired steady level of inflation. That is, the optimal price level this period is the optimal level last period plus some optimal change (which may be zero). The alternative, *rate targeting*, is

$$p^*_t = p_{t-1} + \pi^*, \quad (3)$$

where the current target price level is just the last period's *realized* price level plus the optimal change.

The difference between price-level and inflation rate targeting is the path for the variance of prices. Level targeting implies more volatile short-horizon prices and less volatile long-horizon prices than does rate targeting. To see this, simply note that equation 3 implies that

$$p^*_t = \pi^* t + \sum_{i=0}^{\infty}(p_{t-i} - p^*_{t-i}),$$

which can be a random walk.[6]

The description of the loss function is now complete. It is a function of the parameter vector $\theta = \{\alpha, \beta, h, \pi^*\}$. The values of each of these will depend on the underlying economic structure, that is, tastes and technology. The preference for paths with greater or lesser degrees of variability in output relative to variability in prices, as embodied in the loss function, depends on the fundamental reason that these things are costly. The same is true of the desired steady level of inflation, π^*.

The policymaker's problem cannot be solved without knowledge of the dynamics of output and prices as functions of the policy control variable and the stochastic forcing process driving the economy. These relations, which are taken as constraints in the optimization problem, describe the structure of the economy. For the purposes of the current discussion, I will assume that the central bank policy is carried out using an interest rate, r_t,[7] and that the innovations to the economy come from a series of real and nominal shocks (that is, aggregate demand and aggregate supply shocks), which can be written as ε_t.[8] The reduced form for the evolution of output and prices can then be written as

(4) $$\begin{bmatrix} y_t \\ p_t \end{bmatrix} = A(L) \begin{bmatrix} \varepsilon_t \\ r_t \end{bmatrix},$$

where $A(L)$ is an $(n+1) \times 2$ matrix of (possibly infinite-order) lag polynomials in the lag operator L.[9] The coefficients in $A(L)$ describe a reduced form of the economy. For the moment, I will ignore the fact that $A(L)$ is likely to change when the policy rule changes.[10]

We can now characterize the policymaker's problem as choosing a path for r_t that minimizes the loss (equation 1), with either equation 2 or equation 3 substituted in for p^*, subject to equation 4. The result is a policy rule, which I will write as

(5) $$r_t = \phi(L)\varepsilon_t,$$

where $\phi(L)$ is a (possibly infinite-order) lag polynomial.[11] This path for interest rates as a function of the innovations to the economy (which could be written as differences in the observable quantities) is the policy rule. Significantly, $\phi(L)$ is a function of the parameters θ, as well as the coefficients in $A(L)$ and the covariance matrix of ε, Σ.

I would like to emphasize that the preferences for the evolution of output and price variability, as well as the optimal steady inflation rate π^*, are *inputs* into the policymaker's problem.[12] In practice, I expect that these inputs would be dictated by some legislative or executive body in the government, as they are in some countries (although not in the United States). Given this objective (the loss function) and a model for the evolution of output and prices (the economy), the policymaker chooses a rule that governs the path of the control (the interest rate).[13]

Conceptual Issues and Practical Considerations

The preceding section provides an analytical framework for understanding the policymaker's task, or "problem." In this section, I use this framework to explore several issues relating to policy formulation. Although I leave many questions unanswered, my approach casts new light on some old problems and suggests directions for future research.

I will consider five issues. I begin by exploring the nature of a target. I proceed to a discussion of the practical problems posed by the apparent steepness of the output-inflation variability trade-off and consider how it might influence decisions. This is followed by a general discussion of how uncertainty affects policymaking. Next, I discuss how the nonlinearity created by the fact that the nominal interest rate cannot fall below zero influences the policy rule. Finally, I explore the issue of interest rate smoothing.

Policy Targets

If we accept the view that policy formulation is essentially the solution to the analytical problem of choosing a path for a control variable given a loss function, then how should we interpret the current debate over the proper choice of a policy target, and the advisability of targeting in general? I will explore two ways of addressing the issue of targets. The first is purely technical, and the second has to do with the way policymakers might portray their objective to the public. Technically, the first-order conditions (or Euler equations) to the loss minimization problem I describe above may be interpreted as producing a type of targeting regime. To see this, consider the case examined in detail by Svensson (1996b). He considers pure inflation rate targeting and a loss that is independent of output variation $(\alpha = 1)$. The first-order condition of this problem implies setting the path for expected inflation, $E_t\pi_{t+i}$, as close to the optimal value, π^*, as possible. Svensson refers to this as "inflation forecast targeting."[14]

The only case I can see for intermediate targeting is that it contributes to policy transparency.

This analysis can then be used to justify public statements by policymakers that they are targeting inflation forecasts, as a rhetorical device that substitutes for the more complex and less accessible statements that would be needed to describe their entire procedure.

Ball's (1997) analysis suggests another justification for targets. The argument is that the loss minimization procedure of the type described in the preceding section is too difficult to explain to the population at large (and possibly their elected representatives as well), and so will not lead to policy that is transparent enough to ensure the proper level of accountability.[15] But a pure inflation targeting rule is easy to explain and, more important, easy to understand and monitor. As a result, if the solution to the complex problem can be approximated by a simple rule, there may be substantial virtue in adopting the approximate solution.

A related issue concerns the usefulness of intermediate targets. Over the last half-century or so, many monetary economists have advocated targeting various monetary aggregates. Consider the example of M2.[16] Researchers do not claim to care about M2 for its own sake, nor do they claim that central banks can control it exactly. Therefore, M2 is neither a direct objective nor an instrument. Instead, it is somewhere in between—an *intermediate* target—and the target path would again be akin to the first-order conditions of the optimal control problem.

I find it difficult to make an argument for monetary aggregates as intermediate targets. To see why, consider the case in which the policymaker controls an interest rate and cares about the price level $(\delta = 1)$. To control the objective, the policymaker must know how prices respond to changes in the exogenous environment (the response of p_t to ε_t) and how the objective responds to changes in the instrument. But how does an intermediate target such as M2 help? Clearly, if the relationship between interest rates and M2 and that between M2 and prices are both stable and precisely estimable, then looking at the two relationships separately yields no advantage. In some instances, estimating the impact of interest rates on M2 and the impact of M2 on prices separately might give a more reliable estimate of the product of the two, but such instances would surely be rare. If M2 helps forecast prices, then it will be included in the model. But there is substantial evidence, some of which is in Cecchetti (1995), that reduced-form inflation forecasting relationships are very unstable even if they include M2 or any other potential intermediate target.[17]

As a result, the only case I can see for intermediate targeting is that it contributes to policy transparency. Svensson (1996b) describes an ideal intermediate target that "is highly correlated with the goal, easier to control than the goal, easier to observe by both the central bank and the public than the goal, and transparent so that central bank communication with the public and public understanding and public prediction of monetary policy are facilitated" (pp. 14-5). But since monetary aggregates cannot be closely controlled, are only weakly correlated with both output and inflation over horizons of months or even several years, and have changing definitions that make them difficult to explain, they fail to meet most of Svensson's criteria.

The Output-Inflation Variability Trade-Off

One of the most important practical issues facing policymakers concerns the output-inflation variability trade-off. To measure this trade-off, I turn to some earlier empirical estimates of the impact of central bank policy on output and prices (Cecchetti 1996). In effect, these estimates are the lag polynomials in equation 4 associated with r_t.[18] Chart 1 plots the *impulse response functions*, or dynamic reactions of prices and output to innovations (ε_t's), on the same vertical scale. The most important point to note is that the impact of policy innovations on output is both large and immediate. By contrast, policy affects prices only very slowly, and by much more modest amounts. Furthermore, the precision of the estimates is quite poor.

It is important to keep in mind that standard econometric methods, such as those employed here, assume that parameters are constant over significant historical periods. That is, the vector autoregression (VAR) method used to estimate the response of output and prices to interest rate movements presumes that these reactions are fixed over the 1984-95 sample used in the estimation. Numerous things can cause these relationships to change. As I emphasize in Cecchetti (1995), shifts in central bank policy regimes during the 1970s and 1980s will result in changes in the impulse response functions plotted in Chart 1. In the context of the current discussion, this means that I can reliably measure the output-inflation variability trade-off given the policy regime that was in place over the decade ending in 1995. I cannot reliably estimate the impact of dramatic changes in the policy regime on the trade-off.[19]

Chart 1

RESPONSE OF OUTPUT AND PRICES TO POLICY INNOVATIONS

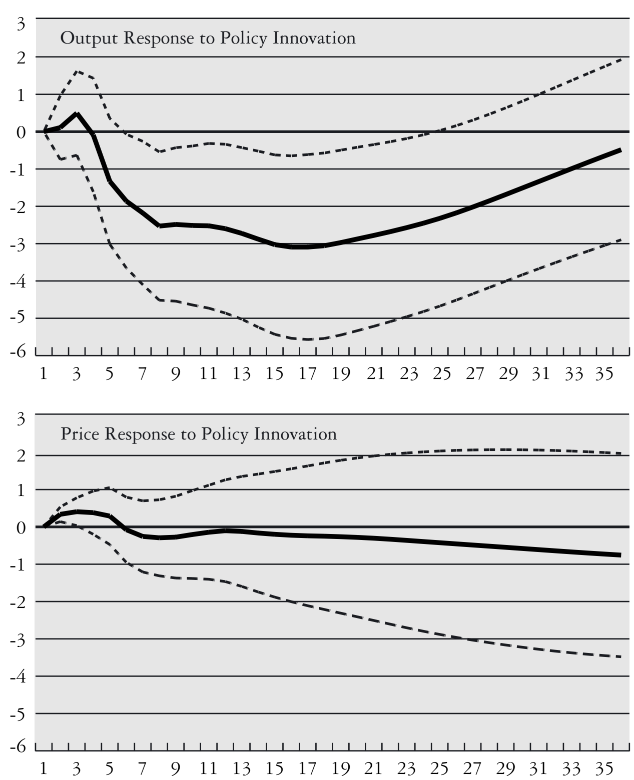

Note: The dotted lines represent standard deviation bands of ±2.

To continue, with the aid of a very simple model, these estimates can be used to give some sense of the shape and slope of the output-inflation variability trade-off. Consider the simple one-period case in which the horizon in the policymaker's loss function (h) is zero, the discount factor (β) is irrelevant, target levels of output (y^*) and prices (p^*) are zero (in logs), and the structure of the economy is such that

(6) $$y_t = \gamma r_t + d_t - s_t, \ \gamma < 0 \quad \text{and}$$

(7) $$p_t = -r_t + d_t + s_t,$$

where d_t and s_t are aggregate demand and aggregate supply shocks. Demand shocks raise both output and prices, while supply shocks move them in opposite directions. I assume that the two types of shocks are uncorrelated and that the variance of the supply shocks is normalized to one, while the variance of the demand shocks is given by σ_d^2.[20] The parameter γ is a measure of the impact of policy innovations on output relative to their impact on prices. The example is meant to represent the medium-horizon impact of policy on the variables of interest. In this simple linear case, the policy rule will be

(8) $$r_t = ad_t + bs_t.$$

Equation 8 implies that

(9) $$\sigma_y^2 = (\gamma a + 1)^2 \sigma_d^2 + (\gamma b - 1)^2 \quad \text{and}$$

(10) $$\sigma_p^2 = (1 - a)^2 \sigma_d^2 + (1 - b)^2.$$

Minimizing the loss function

(11) $$L = \alpha \sigma_p^2 + (1 - \alpha) \sigma_y^2 \quad \text{yields}$$

(12) $$\alpha = \frac{\alpha - \gamma(1 - \alpha)}{\alpha + \gamma^2(1 - \alpha)} \quad \text{and}$$

(13) $$b = \frac{\alpha + \gamma(1 - \alpha)}{\alpha + \gamma^2(1 - \alpha)}.$$

Substituting these into the variance expressions 9 and 10 yields σ_p^2 and σ_y^2 as functions of α, γ, and σ_d^2.

Using the monthly data from Cecchetti (1996), I can now compute an approximate output-inflation variability frontier. From Chart 1, I approximate γ as the mean of the response of output to the mean of the response of prices to an interest rate shock. The result, the average value over a thirty-six-month horizon for the 1984-95 period, is -6.74. The medium-run horizon chosen for these calculations is relevant to policymaking. In a full multiperiod framework, the definition of γ would be more complicated, but its interpretation would remain the same.

Once I determine σ_d^2, the ratio of the variance of demand shocks to the variance of supply shocks, then varying α allows construction of the frontier. Setting σ_d^2 to 0.46 forces the frontier to pass through the value implied by the data (the ratio of output to price variability is approximately 3.72), and normalizing the variance of the detrended log price level in the data to be equal to one gives Chart 2. The "X" marks the value implied by the data. Note that the 1984-95 data suggest that policymakers were operating as if α were approximately 0.93. This is consistent with the importance attached to low and steady inflation over this period.

Chart 2

THE INFLATION-OUTPUT VARIABILITY TRADE-OFF
1984-95

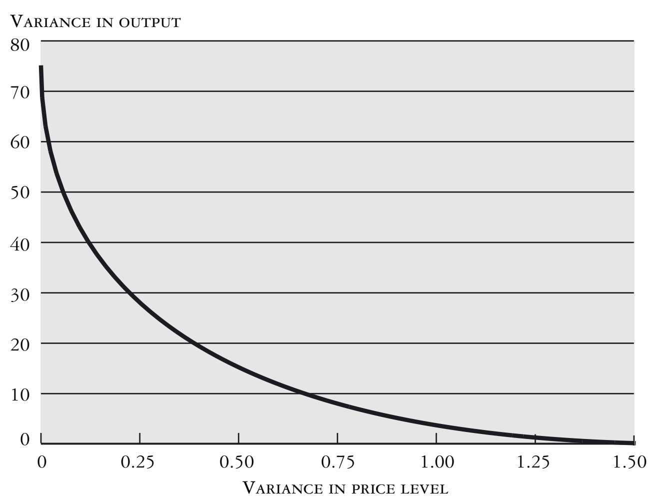

Source: Author's calculations.

Significantly, Chart 2 shows that the trade-off is extremely steep. Reducing inflation variability entirely by setting $\alpha = 1$ creates an extremely high level of variability in real output. In fact, moving from the historically observed point where the ratio of output to inflation variability is 3.72, setting σ_p^2 to zero would increase the variability of output by a factor of more than twenty! By contrast, reducing output variability from 3.72 to zero increases price variability from 1.0 to 1.65. This finding is not a consequence of the simplicity of the example, but rather of the fact that γ is so large. It is straightforward to show that the maximum value of σ_y^2, at $\alpha = 1$, is γ^2 times the maximum value of σ_p^2, at $\alpha = 0$, minus one. That is to say, the points where the line in Chart 2 intersects the x- and y-axes are solely determined by the size of the ratio of policy innovations' impact on output to policy innovations' impact on prices.[21]

Someone who cares about output variability is made substantially worse off by a decision to target the path of the price level.

This result has important implications for the current policy debate. As many central banks move toward some form of price-level or inflation targeting, they are implicitly changing the relative importance of output and inflation variability in their objective function, raising α toward one. From a purely pragmatic point of view, someone who cares about the aggregate price path loses little by allowing α to be less than one. The reverse, however, is emphatically not true. Someone who cares about output variability is made substantially worse off by a decision to target the path of the price level. As a result, when considering policies based on prices alone, policymakers must be very cautious and ask whether they really care so little about output and other real quantities.

Because the estimate of γ plays a crucial role in these conclusions, some comment on its statistical properties is in order. Unfortunately, the estimate is extremely imprecise, with an estimated standard error in excess of 18.[22] This difficulty almost surely stems from the relative stability of inflation during this period and the small and imprecisely estimated response of aggregate prices to policy innovations.[23] The immediate implication is that it is very difficult to be confident of the slope of the variability trade-off. It could be somewhat better, but it could also be substantially worse. A natural reaction to this is to examine the implications of uncertainty for the optimal policy rule.

Uncertainty

How does uncertainty affect policy? Of the numerous types of uncertainty that might influence central bank policy decisions, two forms are examined here: uncertainty about the impact of policy changes (on output and prices) given the model of the economy, and uncertainty about the model itself.[24]

It is straightforward to consider the first of these, which is the sampling error from the estimation of the reaction of prices and output to changes in the policy instrument. In the simple example here, this is just the variance of the estimated γ, which I will call $\sigma^2_{\hat{\gamma}}$. Brainard (1967) originally noted that this type of uncertainty leads to caution in that policy rules imply smaller reactions.[25] In this simple example, inclusion of $\sigma^2_{\hat{\gamma}}$ implies that the policy parameters *a* and *b* become

(14) $$a = \frac{\alpha - \hat{\gamma}(1-\alpha)}{\alpha + (\hat{\gamma}^2 + \sigma^2_{\hat{\gamma}})(1-\alpha)} \quad \text{and}$$

(15) $$b = \frac{\alpha + \hat{\gamma}(1-\alpha)}{\alpha + (\hat{\gamma}^2 + \sigma^2_{\hat{\gamma}})(1-\alpha)} .$$

Reactions to a given shock are now smaller.

In a more realistic, multiperiod context, accounting for parameter uncertainty can be very difficult. Is it likely to be worth the trouble? To get some sense of the impact of parameter uncertainty, I use the results from the previous exercise. If, as I found there, the estimate of γ has a standard error equal to 18.7, then the variance will be 350! The results, plotted in Chart 3, suggest that the impact is huge: the variability frontier shifts out dramatically. Employing the same methods as in Chart 2, I have computed the implied value for σ^2_d so that the inflation-output variability frontier again goes through the point implied by the data. This occurs when σ^2_d equals 0.08 and α equals 0.89, compared with 0.46 and 0.93 in the certainty case. Interestingly, assuming that policymakers have minimized the simple loss function (equation 11) in the presence of uncertainty leads one to conclude that aggregate demand shocks have been substantially less important.

But the real implication of uncertainty is that the frontier is now substantially steeper, and the reaction function parameters *a* and *b* are significantly smaller. In fact, taking account of the changes in both the parameters and the average size of a typical demand or supply shock, I conclude that uncertainty leads to reactions that are on the order of one-twenty-fifth what they were in the certainty case.

What about model uncertainty? There are essentially two problems here. First, past history may not be a reliable guide to the impact of future policy actions, because underlying economic relationships, which policymakers had previously been able to exploit, may change. Such changes could be brought about by policy itself. It is this point, first noted by Lucas (1976), that has driven many macroeconomists to work on dynamic general equilibrium models with well-articulated microeconomic

Chart 3

Impact of Uncertainty on the Variability Trade-off

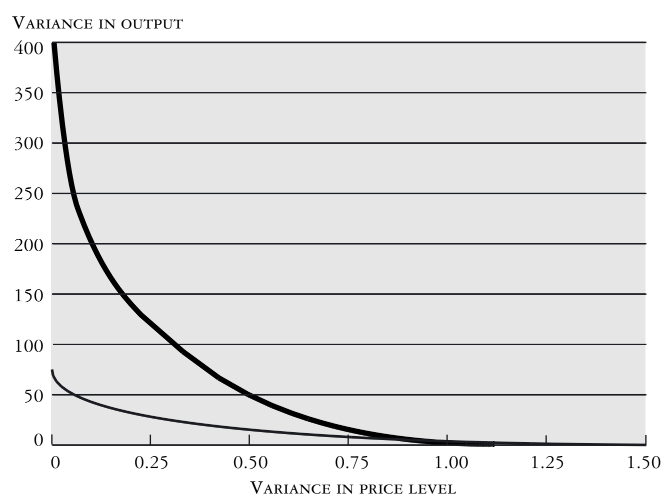

Source: Author's calculations.

foundations. But these efforts are still at too early a stage to be of practical use.

Second, there is little agreement over the true structural model of the economy. McCallum (1997) argues convincingly that, as a result of this lack of consensus, a policy rule should be robust to the possibility that numerous models are correct. In the context of the analytical framework presented in this essay, identifying such a rule would mean exploring the implications of various $A(L)$ s, each of which corresponds to a different model. The object would be to look for a rule that would perform well for a wide range of choices. One method for handling model uncertainty would be to treat it as variance in the estimate of the parameters in $A(L)$.[26] Overall, however, I am forced to conclude that we know very little about how to solve this problem.

The Zero Nominal Interest Rate Floor

What average inflation level should the policymaker target? There are two parts to this question. First, what is the optimal level of inflation, π^*? Second, should policy allow the average realized level of π to deviate from this level?

I argued above that π^* should be dictated to the central bank by social welfare considerations. Quite a bit of work has been done on the possible labor market benefits of modest inflation, suggesting that the optimal level may exceed zero. Most recently, Akerlof, Dickens, and Perry (1996) and Groshen and Schweitzer (1997) consider whether small positive levels of aggregate inflation can facilitate real adjustments in the presence of an aversion to nominal wage declines. But Feldstein (1996) contends that the tax distortions created by inflation reduce the level of output permanently, an argument that suggests π^* may even be negative.[27] Overall, we await further research for the definitive resolution of this issue.

There is one dominant argument for why policymakers might choose to allow average inflation to deviate systematically from the optimal level. The argument, raised in Summers (1991), concerns the case in which π^* is zero, and focuses on the fact that the nominal interest rate cannot fall below zero. In fact, any choice of π^* bounds the real interest rate. Summers goes on to note that in the historical record, the real interest rate (at least ex post) has often been negative. But if central bank policymakers successfully target zero inflation, then the fact that the nominal interest rate cannot be negative means that the real interest rate must always be positive. In essence, this restricts the ability of the policymaker to respond to certain shocks. The control problem as it is described above does not explicitly consider the fact that r_t is bounded at zero. As a result, there will be realizations of ε_t in which the policy rule (equation 5) would imply negative values for the nominal interest rate. One interpretation of Summers' point is that negative nominal interest rates may in some instances be desirable, with the result that mean inflation may deviate from the optimal level in order to allow for a complete response to some larger set of shocks.

In general, the greater the likelihood of a shock driving the desired nominal interest rate below zero, and the higher the loss associated with not being able to react to such a shock, the higher will be the average level of inflation that minimizes the policymaker's loss function.

To see the point, consider the simple model presented in the discussion of the output-inflation variability trade-off. Then, the restriction that $r_t \geq 0$ implies that the loss is minimized for target inflation equal to approximately $0.276\sigma_u$. That is, average inflation will be approximately one-quarter of the standard deviation of the shocks to the price level. More complex forms of the model will have similar properties. In general, the greater the likelihood of a shock driving the desired nominal interest rate below zero, and the higher the loss associated with not being able to react to such a shock, the higher will be the average level of inflation that minimizes the policymaker's loss function.

A similar result would arise when the loss function is asymmetrical. It has been argued that deflation brings potential costs that are distinct from those that come from realized inflation that is less than expected. These costs arise largely because the zero nominal interest rate floor implies that deflation beyond a certain level increases the real interest rate (ex ante and ex post), resulting in a lower steady-state capital stock.[28] This relationship suggests that realized prices below the target may be more costly than equivalent realizations above the target. This would naturally create a positive bias in the policy rule that would result in average inflation exceeding π^*.

To gauge the extent of this problem, I compute the frequency with which the ex post real interest rate has been below zero and below -1.0 percent (see table). Note that the problem is clearly most severe for the United States and France. But for other countries it is relatively modest. In fact, assuming that inflation includes an upward bias of roughly 1 percentage point, the realized real interest rates were negative less than 20 percent of the time in all countries except the United States.

Interest Rate Smoothing

Another important issue for central bank policymakers concerns the desirability of smoothing the changes in the policy instrument. There are two issues here. First is the question whether, following a shock, the optimal response is to have interest rates move immediately up (or down) and then return smoothly down (or up) to the steady-state level, always moving in the same direction following the initial jump. Second, if policymakers intend to change interest rates by some amount, should the entire change occur all at once?

The policy reaction function immediately yields the answer to the first question. Here the presumption must be that $\phi(L)$ is not monotone. That is, it does not imply movements in which interest rates jump initially and then return to the initial level, always moving in the same direction. To see this, consider Chart 4, which plots the optimal reaction of interest rates to an innovation in the aggregate price level implied by the impulse response functions plotted in Chart 1 (for the case where $h = 36$ and $\alpha = 1$).[29] The path is hump-shaped. That is, the optimal response to an innovation is to raise interest rates immediately, continue to raise them gradually, and then lower them slowly. This pattern could be further exaggerated if the loss function included an explicit cost to changing interest rates—a term of the form $k(r_t - r_{t-1})^2$.

The second question is more difficult. If the central bank were to decide that the interest rate should be increased by 100 basis points, should the change be in one large jump or in a series of smaller ones? If policy were sufficiently transparent that everyone knew that the interest rate would ultimately rise 100 basis points, so that the changes would be perfectly anticipated, then it is difficult to see why a series of smaller changes would be

Frequency of Negative Ex Post Real Interest Rate

Country	Date of Initial Observation	Percentage of Observations That Are: Greater Than Zero	Greater Than -1.0
France	Jan. 1970	68	82
Japan	Nov. 1978	82	96
Germany	Jan. 1970	94	99
Italy	Nov. 1979	94	95
United Kingdom	June 1974	77	80
United States	Jan. 1970	69	78

Notes: The interest rate is the rate on three-month Treasury securities, or equivalent. Data are monthly.

Chart 4

Interest Rate Path Following a Shock

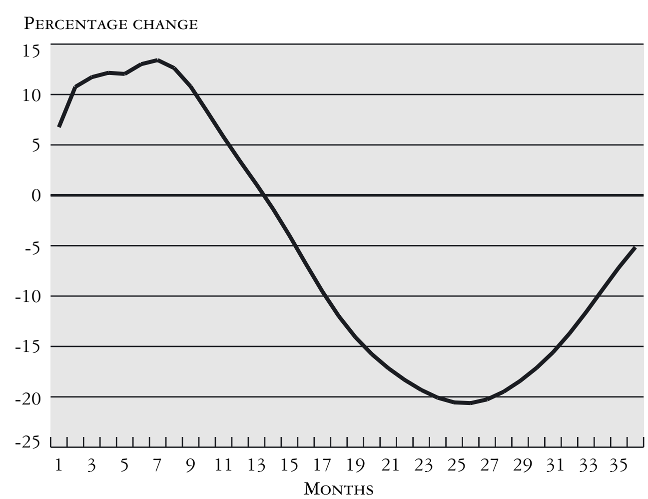

Source: Cecchetti (1996).

preferred over a single one. But often, I suspect, this question is asked with a different intention. In fact, policymakers will start to change interest rates without really knowing what the final results are likely to be. That is, uncertainty about the likely impact of the policy action on the objective will prompt policymakers to make gradual moves so that they can monitor the results—a strategy that may help improve the precision of policymaking.

SHOULD CENTRAL BANKERS FOLLOW RULES?

The entire discussion thus far has been directed at the construction of a rule for central bank policy. But what is our motivation for constructing a set of systematic responses to external events? There are two important reasons to support the adoption of rules by the central bank. The first is the well-known finding that, when policymaking is based on pure discretion rather than rules, the dynamic inconsistency problem leads to high steady inflation. The second reason concerns the importance of policy transparency.

Over fifteen years ago, Barro and Gordon (1983) noted that if a policymaker cannot credibly commit to a zero inflation policy, then even if the policymaker announces that inflation will be zero and all private decisions are based on the assumption that inflation will in fact be zero, it is in the policymaker's interest to renege and induce inflation of some positive amount. The reason for this is that at zero inflation, the value of the increase in output obtained from fooling private agents and creating a transitory increase in output (along a Phillips or Lucas supply curve) more than offsets the cost of the higher inflation, and so the claim of zero inflation in the absence of commitment is not credible. In the language of optimal control, a zero inflation policy is not dynamically consistent.

Since the problem is thought to be most severe when potentially short-sighted legislators are capable of influencing central bank policy directly, the most prominent solution has been to create independent central banks. It is commonly thought, and the data confirm, that policymakers who are more independent are better able to make more credible commitments to low-inflation policy. [30]

As Alan Blinder (1997) has recently pointed out, however, there is a potential conflict between central bank independence and representative democracy. Since one of the crucial elements of a democratic society is that the powerful policymakers are accountable to the people, how can we square these two apparently disparate goals of accountability and independence?

Blinder (1997) and Bernanke and Mishkin (1997) suggest that the solution is policy transparency. They argue that if policymakers announce targets and are forced to explain their actions in relation to these preannounced goals, then there is accountability. Put another way, transparency and accountability are enhanced if the elected officials announce the loss function that the central bankers are charged with minimizing, and if the central bankers in turn demonstrate how they are accomplishing this goal. Researchers have suggested that the publication of the target paths for prices and/or output would serve this purpose. In fact, not only would policymakers become more accountable, but their policies would become more transparent. [31]

Arguments such as these have led to the implementation of explicit targeting regimes in a number of countries. Prominent among these countries are Australia, Canada, Finland, Israel, New Zealand, Spain, Sweden, and the United Kingdom.[32] Because explicit targeting regimes are transparent, they are easily understood. As a result, potential policy actions are less likely to create uncertainty and instability.

LESSONS FOR POLICY FORMULATION

This analysis offers a number of lessons. First, and most important, if a policymaker were to focus on inflation alone, the likely result—in the absence of fundamental changes in the structure of the economy—would be a very high level of real output variation. This finding provides strong support for the very flexible way in which policy targeting is currently carried out around the world.

Consider the example of the countries that have adopted explicit inflation targeting—Australia, Canada, Finland, Israel, New Zealand, Spain, Sweden, and the

United Kingdom, among others. The central banks in most of these countries appear to take into account short-to-medium-run real fluctuations when deciding on their policies. This approach is most evident in the banks' official statements. For example, the central banks in New Zealand, the United Kingdom, and a number of other countries announce target ranges—rather than point targets—for inflation. The Reserve Bank of Australia states that its goal is to have inflation average between 2 and 3 percent *over the business cycle*. By using this wording, the central bank retains the flexibility to stabilize in the face of short-run real shocks. Even countries with explicitly stated inflation targets behave as if they place some weight on output variability in their implicit loss function.

No country has adopted a zero inflation target, or even a range that is centered at zero. In fact, Haldane (1995, p. 8) reports that only New Zealand's target range includes zero at the lower end. This suggests that countries continue to be wary of the possibility of deflation and sensitive to the dangers inherent in bumping against the zero nominal interest rate floor.

The calculations in this essay also underscore the high degree of uncertainty attending the analysis of central bank policy rules. First, I note that the estimated responses of output and inflation to innovations in interest rate policy are extremely imprecise. In other words, policymakers are very unsure about the likely impact of their actions on their objectives. Since I am able to quantify this uncertainty, I can proceed to measure its impact on optimal policymaking. Thus, when I explicitly account for the imprecision of the econometric estimates needed to formulate a rule, I find that the optimal reaction of interest rates to external economic shocks declines by a factor of about twenty-five.

Let me conclude by emphasizing that substantial work remains to be done before we can convincingly articulate a detailed and operational rule for central bank policy. The framework I have proposed requires crucial information on which there is simply no general agreement at this date. What is most needed is a set of stable numerical estimates of the impact of policy actions on output and prices—as well as the ability to estimate the impact of exogenous shocks on the goals of policy.

ENDNOTES

This essay is a revised version of a paper prepared for the Symposium on Operations Research 1997*, held at Freidrich-Shiller University, Jena, Germany, on September 3-5, 1997. The author thanks Paul Bennett, Kenneth Kuttner, Margaret Mary McConnell, Lars Svensson, and Dorothy Sobol for comments. The views stated herein are those of the author and not necessarily those of the Federal Reserve Bank of New York or the Board of Governors of the Federal Reserve System.*

1. Here I paraphrase a comment made by Bennett McCallum at a conference on monetary policy in January 1993.

2. A naive reading of the recent work of Christiano, Eichenbaum, and Evans (1994a, 1994b) surely could lead to such an interpretation.

3. Throughout the discussion in this section, I assume that policymakers can credibly commit to whatever rule they choose. I return to this issue in my discussion of the dynamic consistency problem later in the essay.

4. In some formulations, the loss function includes an additional term in the change in the control variable. That is, changes in interest rates are assumed to be explicitly costly. Inclusion of such a term here adds very little to the analysis.

5. Specifying a process for y^* would be difficult because there is no agreement on several crucial issues. For example, should y^* have a random walk component or be a deterministic trend? Is the growth rate in y^* affected by the volatility of y ?

6. It is possible to nest these two objectives into a more general formulation. Consider a parameter δ representing the relative weight given to price-level targeting and inflation targeting. Then $p_t^* = \delta(\pi^* t) + (1-\delta)(p_{t-1} + \pi^*)$. The percentage of the variance in p explained by its random walk component will be related to δ .

7. The use of an interest rate is not necessary. The control variable could be any quantity that is directly governed by the central bank. For example, the monetary base or some measure of reserves could be used as the control.

8. More specifically, this is a mean zero n-variate stochastic process with finite second moments.

9. Equation 4 is the vector moving-average form. The more common vector autoregression (VAR) is equivalent.

10. This point is emphasized, for example, in Cecchetti (1995).

11. The linear-quadratic structure of the problem described here will give rise to a linear policy rule. In most cases, however, the problem would be structured differently, and the resulting rule would be more complex. For example, if the loss function were nonlinear, or there were some additional constraints on the policymaker's behavior not considered here, then the policy rule would be nonlinear as well.

12. Svensson (1996a) compares inflation and price-level targeting, arguing that one yields better performance than the other under certain economic conditions. Such an exercise relies on a particular view of the costs of inflation that is not explicitly embodied in the loss function (equation 1).

13. Ball (1997) takes a different approach, examining how the adoption of ad hoc rules that are not derived directly from the loss function will affect the loss. For example, after determining the minimized value of the loss L , he then asks how close one can get by adopting a set of arbitrary rules that do not arise from the optimal control problem itself.

14. Svensson (1997) notes that if $\alpha \neq 1$, so that weight is given to output variability in the loss function, the result would be a form of inflation forecast targeting in which the path of the forecast moves gradually back to the optimal level.

15. See the discussion of policy transparency below.

16. For a recent discussion of M2 targeting, see Feldstein and Stock (1994).

17. The procedures of the Deutsche Bundesbank reflect a different view of intermediate targets. As Mishkin and Posen (1997) note, since 1988 the German central bank has targeted growth in M3 in the belief that the demand for German M3 is stable.

18. The methods used to produce these results are described in detail in that earlier work. Briefly, I estimate a four-variable VAR including aggregate prices, commodity prices, industrial production, and the federal funds rate using monthly data from January 1984 to November 1995. Central bank policy innovations are identified and used to estimate the impulse response functions under the assumption that no variable other than policy itself responds to policy shocks immediately.

19. This point is related to the discussion of model uncertainty below.

20. The fact that the coefficients on d and s in equations 7 and 8 are all set to one is a simplification of no consequence for the main point I wish to make. Setting the variance of s to one simply means that the variance of d should be interpreted as the variance of d relative to the variance of s .

21. Cecchetti (1996) considers a substantially more complex case with the same results.

22. There are a number of ways to compute these standard errors. The simplest, called the δ-method, is to note that $\hat{\gamma}$ is a nonlinear function of the estimated parameters of the VAR. A first-order approximation of this function, together with the estimated covariance matrix of the VAR coefficients, can be utilized to compute a variance estimate for γ.

23. It is possible to obtain much more (apparent) precision by computing the average over a longer sample period. For example, using data from 1967 to 1995, the estimate of γ is -2.38, with a standard error of 1.83. But it seems absurd to argue that the reaction function embedded in the VAR is the same now as it was before 1980. This argues strongly for focusing on the estimate from the post-1984 sample.

24. Because of lags in the data, there will also be uncertainty about the current state of the economy. This type of uncertainty has an impact similar to that of parameter uncertainty considered below.

25. Blinder (1997) notes that in a multivariate model, things are not so simple, and the size and sign of covariances will determine whether policymakers exhibit more cautious or less cautious behavior.

26. A simple possibility would be to multiply the estimated covariance matrix of the estimated $A(L)$ by a positive constant.

27. The problem of inflation bias is also relevant here because measured inflation may systematically exceed true inflation. For example, Shapiro and Wilcox (1996) argue that the U.S. consumer price index may overstate inflation by 1 percentage point on average. Such a conclusion suggests that even if π^* is zero, the central bank should attempt to keep measured CPI inflation above zero.

28. See the discussion in Cecchetti (1997).

29. See Cecchetti (1996) for details on this computation.

30. Alesina and Summers (1993) establish this empirically and raise the additional possibility that countries with independent central banks not only have lower steady inflation, but also have less variable output and higher growth. Cukierman et al. (1993) also investigate the impact of central bank independence on the growth rate of output.

31. Mishkin and Posen (1997) argue that policy transparency and explicit targeting were important factors in the granting of operational independence to the Bank of England.

32. See Haldane (1995) for a discussion.

References

Akerlof, George A., William T. Dickens, and George L. Perry. 1996. "The Macroeconomics of Low Inflation." BROOKINGS PAPERS ON ECONOMIC ACTIVITY, no. 1: 1-59.

Alesina, Alberto, and Lawrence H. Summers. 1993. "Central Bank Independence and Macroeconomic Performance: Some Comparative Evidence." JOURNAL OF MONEY, CREDIT, AND BANKING 25 (May): 151-62.

Ball, Laurence M. 1997. "Efficient Monetary Policy Rules." NBER Working Paper no. 5952, March.

Barro, Robert J., and David Gordon. 1983. "A Positive Theory of Monetary Policy in a Natural Rate Model." JOURNAL OF POLITICAL ECONOMY 91 (August): 589-610.

Bernanke, Ben S., and Frederic S. Mishkin. 1997. "Inflation Targeting: A New Framework for Monetary Policy?" NBER Working Paper no. 5893, January.

Blinder, Alan S. 1997. "What Central Bankers Learn from Academics—and Vice Versa." Distinguished Lecture on Economics in Government. JOURNAL OF ECONOMIC PERSPECTIVES 11 (spring): 3-19.

Brainard, William C. 1967. "Uncertainty and the Effectiveness of Policy." AMERICAN ECONOMIC REVIEW 57 (May): 411-25.

Cecchetti, Stephen G. 1995. "Inflation Indicators and Inflation Policy." In B. Bernanke and J. Rotemberg, eds., NBER MACROECONOMICS ANNUAL, 189-219. Cambridge: MIT Press.

————. 1996. "Practical Issues in Monetary Policy Targeting." Federal Reserve Bank of Cleveland ECONOMIC REVIEW 32, no. 1: 2-15.

————. 1997. "Understanding the Great Depression: Lessons for Current Policy." NBER Working Paper no. 6015, April.

Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans. 1994a. "The Effects of Monetary Policy Shocks: Evidence from the Flow of Funds." Federal Reserve Bank of Chicago Working Paper no. 94-2, March.

————. 1994b. "Identification and the Effects of Monetary Policy Shocks." Federal Reserve Bank of Chicago Working Paper no. 94-7, May.

Cukierman, Alex, Pantelis Kalaitzidakis, Lawrence H. Summers, and Steven B. Webb. 1993. "Central Bank Independence, Growth, Investment, and Real Rates." CARNEGIE-ROCHESTER CONFERENCE SERIES ON PUBLIC POLICY 29 (December): 95-140.

Feldstein, Martin. 1996. "The Costs and Benefits of Going from Low Inflation to Price Stability." In C. Romer and D. Romer, eds., REDUCING INFLATION, 123-66. Chicago: University of Chicago Press.

Feldstein, Martin, and James H. Stock. 1994. "The Use of Monetary Aggregates to Target Nominal GDP." In N. Gregory Mankiw, ed., MONETARY POLICY, 7-70. Chicago: University of Chicago Press for NBER.

Groshen, Erica, and Mark Schweitzer. 1997. "Identifying Inflation's Grease and Sand Effects in the Labor Market." NBER Working Paper no. 6061, June.

Haldane, Andrew G. 1995. "Introduction." In A. G. Haldane, ed., TARGETING INFLATION, 1-12. Bank of England, November.

Lucas, Robert E., Jr. 1976. "Econometric Policy Evaluation: A Critique." CARNEGIE-ROCHESTER CONFERENCE SERIES ON PUBLIC POLICY 2: 19-46.

McCallum, Bennett T. 1997. "Issues in the Design of Monetary Policy Rules." NBER Working Paper no. 6016, April.

Mishkin, Frederic S., and Adam Posen. 1997. "Inflation Targeting: Lessons from Four Countries." Federal Reserve Bank of New York ECONOMIC POLICY REVIEW 3, no. 3: 9-110.

Shapiro, Matthew D., and David W. Wilcox. 1996. "Mismeasurement in the Consumer Price Index: An Evaluation." NBER Working Paper no. 5590, May.

Summers, Lawrence H. 1991. "How Should Long-Term Monetary Policy Be Determined?" JOURNAL OF MONEY, CREDIT, AND BANKING 23, no. 3, pt. 2 (August): 625-31.

Svensson, Lars E. O. 1996a. "Price Level Targeting vs. Inflation Targeting: A Free Lunch?" NBER Working Paper no. 5719, August.

————. 1996b. "Inflation Forecast Targeting: Implementing and Monitoring Inflation Targets." NBER Working Paper no. 5797, October.

————. 1997. "Inflation Targeting: Some Extensions." NBER Working Paper no. 5962, March.